arXiv:1512.08829v3 [cs.RO] 7 Nov 2016

Simultaneous Localization And Mapping Without Linearization

Feng Tan, Winfried Lohmiller and Jean-Jacques Slotine

Abstract We apply a combination of linear time varying (LTV) Kalman filtering and nonlinear contraction tools to the problem of simultaneous mapping and localization (SLAM), in a fashion which avoids linearized approximations altogether. By exploiting virtual synthetic measurements, the LTV Kalman observer avoids errors and approximations brought by the linearization process in the EKF SLAM. Furthermore, conditioned on the robot position, the covariances between landmarks are fully decoupled, making the algorithm easily scalable. Contraction analysis is used to establish stability of the algorithm and quantify its convergence rate. We propose four versions based on different combinations of sensor information, ranging from traditional bearing measurements and radial measurements to optical flows and time-to-contact measurements. As shown in simulations, the proposed algorithm is simple and fast, and it can solve SLAM problems in both 2D and 3D scenarios with guaranteed convergence rates in a full nonlinear context.

1 Introduction

Simultaneous localization and mapping (SLAM) is a key problem in mobile robotics research. The Extended Kalman Filtering (EKF) SLAM approach is the earliest and perhaps the most influential SLAM algorithm. It is based on linearizing a nonlinear SLAM model, and then using a Kalman filter to achieve local approximate estimations. However, the linearization process on the originally nonlinear model introduces accumulating errors, which can cause the algorithm to be inconsistent

Feng Tan
Massachusetts Institute of Technology, e-mail: fengtan@mit.edu

Winfried Lohmiller
Massachusetts Institute of Technology e-mail: wslohmil@mit.edu

Jean-Jacques Slotine
Massachusetts Institute of Technology e-mail: jjs@mit.edu

and divergent (Huang and Dissanayake, 2007; Bailey, Nieto, Guivant, Stevens and Nebot, 2006). Such inconsistency is particularly prominent in large-scale problems.

This paper proposes a new approach to the SLAM problem based on creating virtual measurements. Completely free of linearization, this approach yields simpler algorithms and guaranteed convergence rates. The virtual measurements also open up the possibility of exploiting LTV Kalman-filtering and contraction analysis tools in combination.

The proposed algorithm is global and exact, which affords several advantages over existing ones. First, the algorithm is simple and straightforward mathematically, as it exploits purely linear kinematics constraints. Second, following the same LTV Kalman filter framework, the algorithm can adapt to different combinations of sensor information in a very flexible way, and potentially extend to more applications in navigation and machine vision and even contact based localization like Dogar, Hemrajani, Leeds, Kane and Srinivasa (2010) and Javdani, Klingensmith, Bagnell, Pollard and Srinivasa (2013) or SLAM on jointed manipulators like Klingensmith, Srinivasa and Kaess (2016). Third, contraction analysis can be easily used for convergence and consistency analysis of the algorithm, yielding guaranteed global exponential convergence rates. We illustrate the capability of our algorithm in providing accurate estimations in both 2D and 3D settings by applying the proposed framework with different combinations of sensor information, ranging from traditional bearing measurements and range measurements to novel ones such as optical flows and time-to-contact measurements.

Following a brief survey of existing SLAM methods in section II and of basic contraction theory tools in Section III, the algorithm is detailed in section IV for various combinations of sensors. Simulation results are presented in Section V, and concluding remarks are offered in Section VI.

2 A Brief Survey of Existing SLAM Results

In this section we provide a brief introduction on the problem of simultaneous localization and mapping (SLAM). We first review the three most popular categories of SLAM methods: extended Kalman filter SLAM, particle SLAM and graph-based SLAM, and discuss some of their strengths and weaknesses. We then introduce the azimuth model that is used in this paper, along with the kinematics models describing the locomotion of a mobile robot and the landmarks.

Simultaneous localization and mapping is one of key problems in mobile robotics research. SLAM is concerned about accomplishing two tasks simultaneously: mapping an unknown environment with one or multiple mobile robots and localizing the mobile robot/robots. One common model of the environment consists of multiple landmarks such as objects, corners, visual features, salient points, etc. represented by points. And a coordinate vector is used to describe the location of each landmark in 2D or 3D space.

There are three main categories of methods for SLAM: EKF SLAM, graph-based SLAM, and particle filter SLAM. EKF SLAM (Moutarlier and Chatila, 1990; Cheeseman, Smith and Self, 1987; Smith, Self and Cheeseman, 1990; Moutarlier and Chatila, 1989) uses the extended Kalman filter (Jazwinski, 2007; Kalman, 1960), which linearizes and approximates the originally nonlinear problem using the Jacobian of the model to get the system state vector and covariance matrix to be estimated and updated based on the environment measurements.

Graph-based SLAM (Konolige, 2004; Montemerlo and Thrun, 2006; Grisetti, Kummerle, Stachniss and Burgard, 2010; Folkesson and Christensen, 2004; Duckett, Marsland and Shapiro, 2002; Dellaert and Kaess, 2006; Thrun and Montemerlo, 2006) uses graph relationships to model the constraints on estimated states and then uses nonlinear optimization methods to solve the problem. The SLAM problem is modeled as a sparse graph, where the nodes represent the landmarks and each instant pose state, and edge or soft constraint between the nodes corresponds to either a motion or a measurement event. Based on high efficiency optimization methods that are mainly offline and the sparsity of the graph, graphical SLAM methods have the ability to scale to deal with much larger-scale maps.

The particle method for SLAM relies on particle filters (Matthies and Shafer, 1987), which enables easy representation for multimodal distributions since it is a non-parametric representation. The method uses particles representing guesses of true values of the states to approximate the posterior distributions. The first application of such method is introduced in Doucet, De Freitas, Murphy and Russell (2000). The FastSLAM introduced in Montemerlo, Thrun, Koller, Wegbreit et al. (2002) and Montemerlo and Thrun (2007) is one of the most important and famous particle filter SLAM methods. There are also other particle filter SLAM methods such as Del Moral and Guionnet (1999).

However, each of these three methods has weaknesses and limitations:

For the EKF SLAM, the size of the system covariance matrix grows quadratically with the number of features or landmarks, thus heavy computation needs to be carried out in dense landmark environment. Such issue makes it unsuitable for processing large maps. Also, since the linearized Jacobian is formulated using estimated states, it can cause inconsistency and divergence of the algorithm (Huang and Dissanayake, 2007; Bailey et al., 2006).

For graph-based SLAM, because performing the advanced optimization methods can be expensive, they are mostly not online. Moreover, the initialization can have a strong impact on the result.

Lastly, for particle methods in SLAM, a rigorous evaluation in the number of particles required is lacking; the number is often set manually relying on experience or trial and error. Second, the number of particles required increases exponentially with the dimension of the state space. Third, nested loops and extensive re-visits can lead to particles depletion, and make the algorithm fail to achieve a consistent map.

Our method generally falls into the category of Kalman filtering SLAM. By exploiting contraction analysis tools and virtual measurements, our algorithm in effect builds a stable linear time varying(LTV) Kalman filter. Therefore, compared to the EKF SLAM methods, we do not suffer from errors brought by linearization process,

and long term consistency is guaranteed. The math is simple and fast, as we do not need to calculate any Jacobian of the model. And the result we achieve is global, exact and contracting in an exponential favor.

The Azimuth Model of the SLAM Problem in Local Coordinates

Let us first introduce the linear model of SLAM in local coordinates, where we use the azimuth model which measures the azimuth angle in an inertial reference coordinate C_l fixed to the center of the robot and rotates with the robot (Fig.1), as in Lohmiller and Slotine (2013a). The robot is a point of mass with position and attitude.

The actual location of a landmark is described as $\mathbf{x} = (x_1, x_2)^T$ for 2D and $(x_1, x_2, x_3)^T$ for 3D. The measured azimuth angle from the robot is $\theta = \arctan(\frac{x_1}{x_2})$. In 3D there is also the pitch measurement to the landmark $\phi = \arctan(\frac{x_3}{\sqrt{x_1^2+x_2^2}})$. The robot's translational velocity is $\mathbf{u} = (u_1, u_2)^T$ in 2D, and $(u_1, u_2, u_3)^T$ in 3D. Ω is the angular velocity matrix of the robot: in the 2D case, $\Omega = \begin{bmatrix} 0 & -\omega_z \\ \omega_z & 0 \end{bmatrix}$ and in the 3D case, $\Omega = \begin{bmatrix} 0 & -\omega_z & \omega_y \\ \omega_z & 0 & -\omega_x \\ -\omega_y & \omega_x & 0 \end{bmatrix}$. In both cases the matrix Ω is skew-symmetric.

For any landmark $\mathbf{x}_i$ in the inertial coordinate fixed to the robot, the relative motion is:

$$\dot{\mathbf{x}}_i = -\Omega \mathbf{x}_i - \mathbf{u}$$

where both $\mathbf{u}$ and Ω are assumed to be measured accurately, a reasonable assumption in most applications.

If available, the range measurement from the robot to the landmark is $r = \sqrt{x_1^2 + x_2^2}$ in 2D, and $r = \sqrt{x_1^2 + x_2^2 + x_3^2}$ in 3D.

3 Basic Tools in Contraction Theory

Contraction theory (Lohmiller and Slotine, 1998) is a relatively recent dynamic analysis and design tool, which is an exact differential analysis of convergence of complex systems based on the knowledge of the system's linearization (Jacobian) at all points. Contraction theory converts a nonlinear stability problem into an LTV (linear time-varying) first-order stability problem by considering the convergence behavior of neighboring trajectories. While Lyapunov theory may be viewed as a "virtual mechanics" approach to stability analysis, contraction is motivated by a "virtual fluids" point of view. Historically, basic convergence results on contracting systems can be traced back to the numerical analysis literature (Lewis, 1949; Hartman, 1964; Demidovich, 1961).

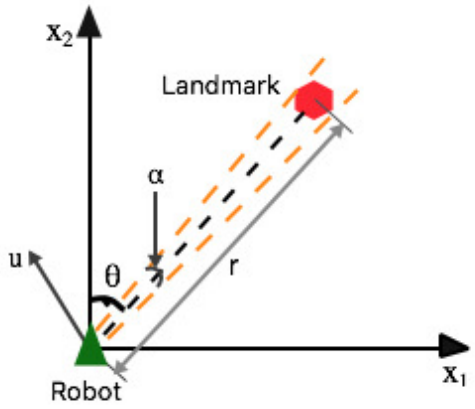

Fig. 1 Various types of measurements in the azimuth model

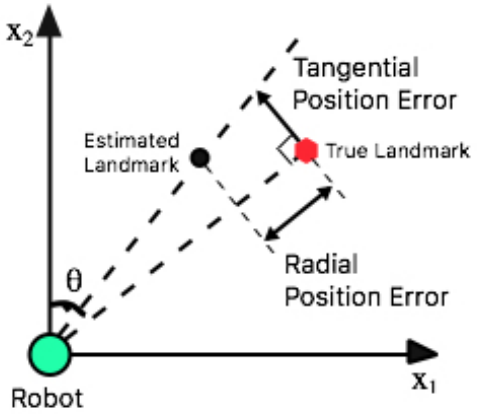

Fig. 2 Position error between estimation and true landmark in Cartesian coordinates

Theorem in Lohmiller and Slotine (1998): *Given the system equations* $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x},t)$*, where* $\mathbf{f}$ *is a differentiable nonlinear complex function of* $\mathbf{x}$ *within* C^n*. If there exists a uniformly positive definite metric* $\mathbf{M}$ *such that*

$$\dot{\mathbf{M}} + \mathbf{M}\frac{\partial \mathbf{f}}{\partial \mathbf{x}} + \frac{\partial \mathbf{f}}{\partial \mathbf{x}}^T \mathbf{M} \leq -\beta_M \mathbf{M}$$

with constant $\beta_M > 0$*, then all system trajectories converge exponentially to a single trajectory, which means contracting, with convergence rate* β_M*.*

Depending on the application, the metric can be found trivially (identity or rescaling of states), or obtained from physics (say, based on the inertia tensor in a mechanical system as e.g. in Lohmiller and Slotine (2008, 2013b)). The reader is referred to Lohmiller and Slotine (1998) for a discussion of basic features in contraction theory.

4 Landmark Navigation and LTV Kalman Filter SLAM

In this section we illustrate the use of both LTV Kalman filter and contraction tools on the problem of navigation with visual measurements, an application often referred to as the landmark (or lighthouse) problem, and a key component of simultaneous localization and mapping (SLAM).

The main issues for EKF SLAM lie in the linearization and the inconsistency caused by the approximation. Our approach to solve the SLAM problem in general follows the paradigms of LTV Kalman filter. And contraction analysis adds to the global and exact solution with stability assurance because of the exponential convergence rate.

We present the results of an exact LTV Kalman observer based on the Riccati dynamics, which describes the Hessian of a Hamiltonian p.d.e. (Lohmiller and Slotine, 2013a). A rotation term similar to that of Grave and Tang (2015) in the context of perspective vision systems is also included.

4.1 LTV Kalman filter SLAM using virtual measurements in local coordinates

A standard extended Kalman Filter design (Bryson, 1975) would start with the available nonlinear measurements, for example in 2D (Fig.1)

$$\theta = \arctan\left(\frac{x_1}{x_2}\right) \qquad \text{and/or} \qquad r = \sqrt{x_1^2 + x_2^2}$$

and then linearizes these measurements using the estimated Jacobian, leading to a locally stable observer. Intuitively, the starting point of our algorithm is the simple remark that the above relations can be equivalently written in Cartesian coordinates as

$$\mathbf{h}\mathbf{x} = 0 \qquad \text{and/or} \qquad \mathbf{h}^*\mathbf{x} = r$$

where for 2D scenarios

$$\mathbf{h} = (\cos\theta, -\sin\theta) \qquad\qquad \mathbf{h}^* = (\sin\theta, \cos\theta)$$

and for 3D scenarios:

$$\mathbf{h} = \begin{pmatrix} \cos\theta & -\sin\theta & 0 \\ -\sin\phi\sin\theta & -\sin\phi\cos\theta & \cos\phi \end{pmatrix}$$

$$\mathbf{h}^* = (\cos\phi\sin\theta, \cos\phi\cos\theta, \sin\phi)$$

Indeed in the azimuth model, one has in 2D

$$\mathbf{x} = (x_1, x_2)^T = (r\sin\theta, r\cos\theta)^T$$

and in 3D

$$\mathbf{x} = (x_1, x_2)^T = (r\cos\phi\sin\theta, r\cos\phi\cos\theta, r\sin\phi)^T$$

Thus, instead of directly comparing the measurement and the observation of θ and r, we choose to have feedbacks on the tangential and radial Cartesian position errors between estimated and true landmark positions using simple geometrical transformation. So instead of the nonlinear observations, we can have the linear substitutes using virtual measurements and further exploit these exact linear time-varying expressions to achieve a globally stable observer design. This simple philosophy can be easily extended to a variety of SLAM contexts with different measurement inputs, especially in the visual SLAM field.

All our propositions in the following cases have the same continuous LTV Kalman filter structure using the implicit measurements

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{v}(t)$$

where $\mathbf{y}$ is the measurement/observation vector in Cartesian coordinates, which includes both actual and virtual measurements; $\mathbf{H}$ is the observation model matrix, consisted of state-independent measurement vectors such as $\mathbf{h}$ and $\mathbf{h}^*$; $\mathbf{v}(t)$ is a zero-mean white noise also in Cartesian coordinates with the covariance $\mathbf{R}$.

The filter consists of two differential equations, one for the state estimations and one for the covariance updates:

$$\dot{\hat{\mathbf{x}}} = -\mathbf{u} - \Omega\hat{\mathbf{x}} + \mathbf{K}(\mathbf{y} - \mathbf{H}\hat{\mathbf{x}})$$

$$\dot{\mathbf{P}} = \mathbf{Q} - \mathbf{P}\mathbf{H}^T\mathbf{R}^{-1}\mathbf{H}\mathbf{P} - \Omega\mathbf{P} + \mathbf{P}\Omega$$

with the Kalman gain $\mathbf{K}$ given by

$$\mathbf{K} = \mathbf{P}\mathbf{H}^T\mathbf{R}^{-1}$$

and

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{v}(t)$$

$$\mathbf{Q} = cov(\mathbf{w})$$

where $\mathbf{w}(t)$ is a zero-mean white noise included in $\mathbf{u}$ (due e.g. to motion measurement inaccuracy, or rough or slippery terrain).

The definition of $\mathbf{y}$ and $\mathbf{H}$ varies according to the type of measurement available, as we now discuss.

Case I: bearing measurement only

This original version of bearing-only SLAM was presented in Lohmiller and Slotine (2013a) where:

$$y = 0 \qquad \mathbf{H} = \mathbf{h}$$

Geometrically, the virtual measurement error term $\mathbf{h}\hat{\mathbf{x}}$ corresponds to rewriting an angular error as a tangential position error between estimated and true landmark positions. As the vehicle moves, at any instant the system contracts exponentially in the tangential direction if $\mathbf{R}^{-1} > 0$, and it is indifferent along the unmeasured radial direction.

Case II: bearing with range measurement

If we have both bearing measurement θ and ϕ and range measurement r, a new constraint would be $y_2 = r = \mathbf{h}^*\mathbf{x}$.

So that for both 2D and 3D the virtual measurement is:

$$\mathbf{y} = \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} 0 \\ r \end{bmatrix}$$

$$\mathbf{H} = \begin{bmatrix} \mathbf{h} \\ \mathbf{h}* \end{bmatrix}$$

Case III: bearing with independent $\dot{\theta}$ information

In this case we utilize $\dot{\theta}$ as additional information. $\dot{\theta}$ is the measured relative angular velocity from the robot to the landmark and we also have $\dot{\phi}$ in the 3D case. Independent $\dot{\theta}$ measurement could be achieved either computationally based on θ or through optical flow algorithms on visual sensors. We propose here that $\dot{\theta}$ gives us an additional dimension of information that helps the LTV Kalman filter with radial contraction. The additional constraint or observation we get is based on the relationship

$$range \times angular\ velocity = tangential\ velocity$$

where in our case $\mathbf{h}^*\hat{\mathbf{x}}$ is the length of vector $\hat{\mathbf{x}}$ projected along azimuth direction to represent the estimated range.

So if the estimation is precise, $\dot{\theta}\mathbf{h}^*\hat{\mathbf{x}} + \mathbf{h}\Omega\hat{\mathbf{x}}$ should equal to $-\mathbf{hu}$, which is the relative velocity projected along the tangential direction.

In this case, $y_1 = \mathbf{hx} = 0$ is the constraint on bearing measurement and $y_3 = (\dot{\theta}\mathbf{h}^* + \mathbf{h}\Omega)\mathbf{x} = -\mathbf{hu}$ in 2D or $\mathbf{y_3} = (\begin{bmatrix} \dot{\theta} \begin{bmatrix} \sin\theta & \cos\theta & 0 \end{bmatrix} \\ \dot{\phi}\mathbf{h}^* \end{bmatrix} + \mathbf{h}\Omega)\mathbf{x} = -\mathbf{hu}$ in 3D is the constraint about relative angular velocity, radial distance and the tangential velocity.

The same constraint could be derived similarly from:

$$\frac{d}{dt}\mathbf{hx} = 0 \Rightarrow \dot{\mathbf{h}}\mathbf{x} + \mathbf{h}\dot{\mathbf{x}} = 0$$

which means

$$\dot{\mathbf{h}}\mathbf{x} + \mathbf{h}(\Omega\mathbf{x} + \mathbf{u}) = 0$$

where $\dot{\mathbf{h}} = \dot{\theta}\mathbf{h}^*$ in 2D and $\dot{\mathbf{h}} = \begin{bmatrix} \dot{\theta}\begin{bmatrix}\sin\theta & \cos\theta & 0\end{bmatrix} \\ \dot{\phi}\mathbf{h}^* \end{bmatrix}$ in 3D. Such derivation achieves the same result as the constraint we proposed earlier about tangential velocity. So the virtual measurement consists of two parts:

$$\mathbf{y} = \begin{bmatrix} y_1 \\ y_3 \end{bmatrix} = \begin{bmatrix} 0 \\ -\mathbf{hu} \end{bmatrix}$$

with the observation model: (2D) $H = \begin{bmatrix} \mathbf{h} \\ \dot{\theta}\mathbf{h}^* + \mathbf{h}\Omega \end{bmatrix}$ (3D) $H = \begin{bmatrix} \mathbf{h} \\ \begin{bmatrix} \dot{\theta}\begin{bmatrix}\sin\theta & \cos\theta & 0\end{bmatrix} \\ \dot{\phi}\mathbf{h}^* \end{bmatrix} + \mathbf{h}\Omega \end{bmatrix}$

Case IV: bearing with time to contact measurement τ

In this case we utilize the "time to contact" measurement as additional information. Time-to-contact (Horn, Fang and Masaki, 2007; Browning, 2012; Clady, Clercq, Ieng, Houseini, Randazzo, Natale, Bartolozzi and Benosman, 2015) measurement provides an estimation of time to reach the landmark, which could suggest the radial distance to the landmark based on local velocity information. This is one popular measurement for sailing and also utilized by animals and insects. For a robot, the "time to contact" measurement could be potentially achieved by optical flows algorithms, direct gradient based methods like Horn et al. (2007), or some novel sensors specifically developed for that purpose.

$$\tau = |\frac{\alpha}{\dot{\alpha}}| \approx |\frac{r}{\dot{r}}|$$

As shown in Fig.1, we can get the measurement $\tau = |\frac{\alpha}{\dot{\alpha}}|$, where α is an small angle measured between two feature points, edges on a single distant landmark for example. In our case, we use the angle between two edges of the cylinder landmark so that $\alpha \approx arctan(\frac{d}{r})$, where d is the diameter of the cylinder landmark and r is the distance from the robot to the landmark. Thus in this case besides the bearing constraint $y_1 = \mathbf{hx} = 0$, we propose a novel constraint y_4 utilizing the "time to contact" τ.

As we know $r = \mathbf{h}^*\mathbf{x}$ so that

$$\dot{r} = -\mathbf{h}^*\mathbf{u} - \mathbf{h}^*\Omega\mathbf{x} + \dot{\mathbf{h}}^*\mathbf{x}$$

Since $\mathbf{h}^*$ is the unit vector with the same direction of $\mathbf{x}$, both $\mathbf{h}^*\Omega\mathbf{x}$ and $\dot{\mathbf{h}}^*\mathbf{x}$ equal to 0, so simply $\dot{r} = -\mathbf{h}^*\mathbf{u}$, and

$$\tau = |\frac{r}{\dot{r}}| = \frac{\mathbf{h}^*\mathbf{x}}{|-\mathbf{h}^*\mathbf{u}|}$$

which means: $|\tau\mathbf{h}^*\mathbf{u}| \approx \mathbf{h}^*\mathbf{x}$ so we can have $y_4 = |\tau\mathbf{h}^*\mathbf{u}|$

$$\mathbf{y} = \begin{bmatrix} y_1 \\ y_4 \end{bmatrix} = \begin{bmatrix} 0 \\ |\tau\mathbf{h}^*\mathbf{u}| \end{bmatrix}$$

$$\mathbf{H} = \begin{bmatrix} \mathbf{h} \\ \mathbf{h}^* \end{bmatrix}$$

So that $\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{v}$, and it is applicable to both 2D and 3D cases. One thing to notice is that the time-to-contact measurement is an approximation. Also when $\mathbf{uh} \approx 0$, τ would be reaching infinity, which reduces the reliability of the algorithm near that region.

Case V: range measurement only

If the robot has no bearing information, it may still perform SLAM if range measurements and their time-derivatives are available. Since

$$r^2 = \mathbf{x}^\mathbf{T}\mathbf{x} \quad \Rightarrow \quad \frac{d}{dt}r^2 = \frac{d}{dt}\mathbf{x}^\mathbf{T}\mathbf{x} \quad \Rightarrow \quad r\dot{r} = -\mathbf{u}^\mathbf{T}\mathbf{x}$$

measurements of both r and $\dot{r}$ (e.g. from a Doppler sensor) can be used in the LTV Kalman filter framework, in which case

$$y = r\dot{r} = \mathbf{H}\mathbf{x}$$

$$\mathbf{H} = \mathbf{u}^\mathbf{T}$$

in both 2D and 3D.

Remark I

The estimated landmark positions are based on the azimuth model in the inertial coordinate system fixed to the robot. Thus, the positions of the landmarks are positions relative to the robot rather than global locations. Denoting the states of the visible landmarks by $\mathbf{x}_{il}$, and corresponding measurements θ_i and r_i, each with independent covariance matrix P_i.

$$\dot{\hat{\mathbf{x}}}_{il} = -\mathbf{u} - \Omega\hat{\mathbf{x}}_{il} + \mathbf{K}(\mathbf{y} - \mathbf{H}_{il}\hat{\mathbf{x}}_{il})$$

$$\dot{\mathbf{P}}_i = \mathbf{Q} - \mathbf{P}_i\mathbf{H}_{il}^T\mathbf{R}^{-1}\mathbf{H}_{il}\mathbf{P}_i - \Omega\mathbf{P}_i + \mathbf{P}_i\Omega$$

Since the relative landmark positions are conditioned on the local inertial coordinates of the robot, covariances on each pairs of landmarks are fully decoupled, which shrinks the covariance matrix down to the dimension of single landmarks coordinates. And complexity in all local cases scales linearly with the number of landmarks.

Remark II

Our cases proposed above provide suggestions to exploit information from different sensor measurements from various instruments under the same framework of LTV Kalman filter. For example, a Doppler radar or sonar measures quite good r and $\dot{r}$; a camera measures the bearings ϕ and ψ, while combined with optical flows and other algorithms, it can also measure $\dot{\phi}$, $\dot{\psi}$ and τ; and a lidar measures good r and ϕ and ψ.

4.2 LTV Kalman filter in 2D global coordinates

Before introducing the LTV Kalman filter in 2D global coordinates, we first analyze the Kalman filter in an intermediate local coordinate system C_L.

4.2.1 LTV Kalman filter in 2D rotation only coordinates

The local coordinate C_L is fixed to the origin of the robot at $t = 0$, and has the same attitude as the robot, which means C_L has no translation movement but rotates as same as the robot. Here we use the bearing only case (Case I) for example to explain the structure, it can be easily extended to other measurement models. Recall:

$$\theta = arctan(\frac{x_1}{x_2})$$
$$\mathbf{h} = (\cos\theta, -\sin\theta)$$
$$\Omega = \begin{bmatrix} 0 & -\omega \\ \omega & 0 \end{bmatrix}$$

We name the local coordinates in C_L for each landmark's position $\mathbf{x}_{iL}$ and the vehicle's position $\mathbf{x}_{vL}$. For the bearing only case we have the linear constraint

$$\mathbf{h}(\mathbf{x}_{iL} - \mathbf{x}_{vL}) = 0$$

For the kinematics of the system we have both linear velocity and angular velocity on the vehicle and angular velocity alone on the landmarks:

$$\dot{\mathbf{x}}_{vL} = \mathbf{u} + \Omega\mathbf{x}_{vL}$$

$$\dot{\mathbf{x}}_{iL} = \Omega\mathbf{x}_{iL}$$

So similar to the LTV Kalman filter proposed in C_l before, we can use another LTV Kalman system updated as:

$$\begin{bmatrix} \dot{\hat{\mathbf{x}}}_{1L} \\ \dot{\hat{\mathbf{x}}}_{2L} \\ \vdots \\ \dot{\hat{\mathbf{x}}}_{nL} \\ \dot{\hat{\mathbf{x}}}_{vL} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ \mathbf{u} \end{bmatrix} + diag(\Omega) \begin{bmatrix} \hat{\mathbf{x}}_{1L} \\ \hat{\mathbf{x}}_{2L} \\ \vdots \\ \hat{\mathbf{x}}_{nL} \\ \hat{\mathbf{x}}_{vL} \end{bmatrix} - \mathbf{P}\mathbf{H}^T\mathbf{R}^{-1}\mathbf{H} \begin{bmatrix} \hat{\mathbf{x}}_{1L} \\ \hat{\mathbf{x}}_{2L} \\ \vdots \\ \hat{\mathbf{x}}_{nL} \\ \hat{\mathbf{x}}_{vL} \end{bmatrix}$$

Here $\mathbf{H} = \begin{bmatrix} \mathbf{h}_1 & 0 & \cdots & 0 & -\mathbf{h}_1 \\ 0 & \mathbf{h}_2 & \cdots & 0 & -\mathbf{h}_2 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & \mathbf{h}_n & -\mathbf{h}_n \end{bmatrix}$, and the covariance matrix updates as:

$$\dot{\mathbf{P}} = \mathbf{Q} - \mathbf{P}\mathbf{H}^T\mathbf{R}^{-1}\mathbf{H}\mathbf{P} + diag(\Omega)\mathbf{P} - \mathbf{P}diag(\Omega)$$

This system is contracting while free of attitude of the vehicle (heading angle β in 2D case). And since the true positions of landmarks and vehicle are particular solutions to this contracting system, all estimated trajectories are guaranteed to converge to the true trajectory exponentially.

4.2.2 LTV Kalman filter in 2D global coordinates

Now we look at the problem of LTV Kalman filter in 2D global coordinates. We call the coordinate system C_G, which is fixed at the starting point of the robot. Positions of landmarks in global coordinates are $\mathbf{x}_{iG}$ and position of the vehicle is $\mathbf{x}_{vG}$. So a transformation matrix from coordinate system C_G to C_l is simply a rotation matrix $\mathbf{T}(\beta)$. It means $\mathbf{x}_{iL} = \mathbf{T}(\beta)\mathbf{x}_{iG}$ and $\mathbf{x}_{vL} = \mathbf{T}(\beta)\mathbf{x}_{vG}$. Substituting these transformations into the LTV Kalman filter proposed in 4.2.1, we can have a LTV Kalman filter in global coordinates describing the same model using virtual measurements as:

$$\begin{bmatrix} \dot{\hat{\mathbf{x}}}_{1G} \\ \dot{\hat{\mathbf{x}}}_{2G} \\ \vdots \\ \dot{\hat{\mathbf{x}}}_{nG} \\ \dot{\hat{\mathbf{x}}}_{vG} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ u\cos\hat{\beta} \\ u\sin\hat{\beta} \end{bmatrix} - \mathbf{P}\mathbf{H}_G^T\mathbf{R}^{-1}\mathbf{H}_G \begin{bmatrix} \hat{\mathbf{x}}_{1L} \\ \hat{\mathbf{x}}_{2L} \\ \vdots \\ \hat{\mathbf{x}}_{nL} \\ \hat{\mathbf{x}}_{vL} \end{bmatrix}$$

Where

$$\mathbf{H}_G = \begin{bmatrix} \mathbf{h}_1\mathbf{T}(\hat{\beta}) & 0 & \cdots & 0 & -\mathbf{h}_1\mathbf{T}(\hat{\beta}) \\ 0 & \mathbf{h}_2\mathbf{T}(\hat{\beta}) & \cdots & 0 & -\mathbf{h}_2\mathbf{T}(\hat{\beta}) \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & \mathbf{h}_n\mathbf{T}(\hat{\beta}) & -\mathbf{h}_n\mathbf{T}(\hat{\beta}) \end{bmatrix}$$

In this LTV Kalman filter, we extract the attitude of the robot from of the state vector and treat it as a component in the virtual measurement, which helps eliminate the nonlinearity in the model. The value of β is determined based on the model by tracking a desired value β_d,

$$\dot{\hat{\beta}} = \omega + (\beta_d - \hat{\beta})$$

where β_d minimizes the quadratic residue error,

$$\beta_d = argmin_{\beta_d \in [-\pi,\pi]} \mathbf{x}^T\mathbf{H}^T\mathbf{H}\mathbf{x}$$

In the 2D case,

$$\beta_d = \frac{1}{2}\arctan\frac{\sum_i[-\sin 2\theta_i(x_{iG1}^2 - x_{iG2}^2) + 2\cos 2\theta_i x_{iG1}x_{iG2}]}{\sum_i[-\cos 2\theta_i(x_{iG1}^2 - x_{iG2}^2) + 2x_{iG1}x_{iG2}\sin 2\theta_i]}$$

when at every instant the sums are taken over the visible landmarks. Since this LTV Kalman filter is obtained simply by a coordinate transformation from the Kalman filter proposed in 4.2.1, which is contracting regardless of the attitude states, this LTV Kalman filter in the global coordinate system is contracting as well exponentially towards the true states.

Because all cases analyzed in previous sections using different combinations of sensor information follow the same framework, they can all be treated in global coordinates as same as the bearing only case. In each of these cases $\mathbf{H}_G = \mathbf{H}_L diag(\mathbf{T}(\hat{\beta}))$, with a different β_d minimizing the residue error $(\mathbf{y}-\mathbf{H}\mathbf{x})^T(\mathbf{y}-\mathbf{H}\mathbf{x})$ in each case.

Note that for estimating the positions of landmarks and the vehicle in global coordinates, we are actually utilizing a full state Kalman filter, so that, computationally, the proposed LTV Kalman filter takes as much computation as traditional EKF methods. However, our LTV Kalman filter is linear, global and exact. Furthermore, it uses a common framework to solve problems involving different combinations of sensor information, and contracts exponentially to the true states.

Remark I: Nonlinearity in vehicle kinematics

When traditional EKF SLAM methods are applied to ground vehicles, another nonlinearity arises from the vehicle kinematics. This is easily incorporated in our model.

The vehicle motion can be modeled as

$$\dot{x}_{v1} = u\ \cos\beta$$

$$\dot{x}_{v2} = u\ \sin\beta$$

$$\dot{\beta} = \ \omega\ = \frac{u}{L}\ \tan\theta_s$$

where u is the linear velocity, L is the distance between the front and rear axles and θ_s is the steering angle. As the heading angle β is an independent input generated by an upstream level of the filter dynamics, the kinematics of the vehicle remains linear time-varying.

$$\frac{d}{dt}\begin{bmatrix} x_{v1} \\ x_{v2} \end{bmatrix} = \begin{bmatrix} u\ \cos\beta \\ u\ \sin\beta \end{bmatrix}$$

Remark II: 3D capability

It is obvious that our proposals in all cases in local coordinates have full capability of dealing with 6DOF problem. That means when based on the local coordinate system fixed on the robot, or based on a local coordinate system fixed at the origin rotating with the robot, we can deal with 3D problems very easily. Such scenarios are more popular for the flying vehicles to map the surrounding environment rather than large scale long history global mapping. For large scale global mapping, LTV Kalman filter proposed in 4.2.2 has been fully proved to work in 2D applications. In 3D(6DOF) scenarios, the LTV Kalman filter by itself still works with no question,

it is just that we may need some other nonlinear optimization methods to find the desired attitude states for the estimated states to track to minimize the residue error. And further we put the estimated value of the attitudes into the LTV Kalman filter as measurements as same as the 2D case.

Remark III: Second-order vehicle dynamics

The algorithm can be easily extended if instead of having direct velocity measurement, the vehicle dynamics model is second-order,

$$\ddot{\mathbf{x}}_{vG} = \mathbf{u}$$

where $\mathbf{u}$ is now the translational *acceleration* instead of the translational velocity, and the state vector is augmented by the linear velocity vector of the vehicle. Cases I to IV extend straightforwardly since the constraints

$$\mathbf{h}_i \mathbf{T}(\mathbf{x}_{iG} - \mathbf{x}_{vG}) = 0$$

$$\mathbf{h}_i^* \mathbf{T}(\mathbf{x}_{iG} - \mathbf{x}_{vG}) = r_i$$

$$(\dot{\theta}_i \mathbf{h}_i^* + \mathbf{h}_i \Omega)\mathbf{T}(\mathbf{x}_{iG} - \mathbf{x}_{vG}) = -\mathbf{h}_i \dot{\mathbf{x}}_{vG}$$

$$|\tau \mathbf{h}_i^* \dot{\mathbf{x}}_{vG}| = \mathbf{h}_i^* \mathbf{T}(\mathbf{x}_{iG} - \mathbf{x}_{vG})$$

remain linear. Only case V (range only) does not, as it relies on the product of the position and velocity of the vehicle in the model.

4.3 Contraction analysis for the LTV Kalman filter

Since all our cases follow the same LTV Kalman filter structure, we can analyze the contraction property in general for all cases at the same time. The LTV Kalman filter system we proposed previously contracts according to Section 3, with metric $\mathbf{M}_i = \mathbf{P}_i^{-1}$, as analyzed in Lohmiller and Slotine (2013a):

$$\frac{\partial \mathbf{f}_i^T}{\partial \mathbf{x}_i}\mathbf{M}_i + \mathbf{M}_i \frac{\partial \mathbf{f}_i}{\partial \mathbf{x}_i} + \dot{\mathbf{M}}_i = -\mathbf{M}_i \mathbf{Q}_i \mathbf{M}_i - \mathbf{H}_{il}^T \mathbf{R}^{-1} \mathbf{H}_{il}$$

The above leads to the global exponential Kalman observer of landmarks (lighthouses) around a vehicle. Hence for any initial value, our estimation will converge to the trajectory of the true landmarks positions exponentially. It gives stability proof to the proposed LTV Kalman filter and boundedness of $\mathbf{M}$ is given with the observability grammian. However, LTV Kalman cannot compute the convergence rates explicitly, because the convergence rate is given by the eigenvalues of $-\mathbf{M}_i \mathbf{Q}_i \mathbf{M}_i - \mathbf{H}_{il}^T \mathbf{R}^{-1} \mathbf{H}_{il}$ which is related to $\mathbf{M}$.

This system is contracting with metric $\mathbf{M} = \mathbf{P}^{-1}$ in global coordinates also. Since the true locations of landmarks and path of the vehicle are particular solutions to the system, all trajectories of the state vectors would converge exponentially to the truth.

4.4 Noise analysis

A basic assumption for the Kalman filter is that the noise signal $\mathbf{v}(t) = \mathbf{y} - \mathbf{H}\mathbf{x}$ is zero-mean. Since the actual measurements obtained from a robot are θ, ϕ, $\dot{\theta}$, $\dot{\phi}$, r, and τ, we need to verify that the mean of noise remains zero after incorporating them into the virtual measurements. Similarly, the variance estimates in the initial noise model have to be ported to the new variables. The general philosophy of this paper is that typically the noise models themselves are somewhat coarse estimates, so that this translation of estimated noise variances to the new variables can be approximate without much practical loss of performance. Recall also that the LTV Kalman filter is the optimal least-squares LTV filter given the means and variances of the driving and measurement noise processes, regardless of the noise distribution. In addition, the precision on $\mathbf{Q}$ and $\mathbf{R}$ does not affect the filters's stability and convergence rates, but only its optimality.

For Case V, since $\mathbf{H} = \mathbf{u}^{\mathbf{T}}$ and radius r are directly measured independently, there would be no mean-shift for the noise $v = \mathbf{u}^{\mathbf{T}}\mathbf{x} + r\dot{r}$, which means $E[v] = 0$ in both 2D and 3D cases.

For all the other cases, assume the bearing angle θ we measure comes with a zero-mean white Gaussian noise $w \sim N(0, \sigma_\theta^2)$, and ϕ with a zero-mean white Gaussian noise $N(0, \sigma_\phi^2)$ then

$$E[\cos(\theta + w)] = e^{-\frac{\sigma_\theta^2}{2}} \cos\theta$$

$$E[\sin(\theta + w)] = e^{-\frac{\sigma_\theta^2}{2}} \sin\theta$$

Combined with the geometry:
(2D)

$$x_1 = r\sin\theta$$

$$x_2 = r\cos\theta$$

(3D)

$$x_1 = r\sin\theta\sin\phi$$

$$x_2 = r\cos\theta\sin\phi$$

$$x_3 = r\sin\phi$$

For our virtual measurement $y = \mathbf{h}\mathbf{x} + v$
where $\mathbf{h} = [\cos\theta, -\sin\theta]$, the noise

$$v = 0 - \mathbf{h}\mathbf{x} = -(x_1\cos\theta - x_2\sin\theta)$$

so that the mean of the noise

$$E[v] = -e^{-\frac{\sigma_\theta^2}{2}}(x_1\cos\theta - x_2\sin\theta) = 0$$

which means there is no bias in this case.

In the 3D case

$$\mathbf{h} = \begin{pmatrix} \cos\theta & -sin\theta & 0 \\ -\sin\phi\sin\theta & -\sin\phi\cos\theta & \cos\phi \end{pmatrix}$$

$$\begin{aligned} E[\mathbf{v}] &= \begin{pmatrix} -e^{-\frac{\sigma_\theta^2}{2}}(x_1\cos\theta - x_2\sin\theta) \\ r(e^{-\frac{\sigma_\phi^2}{2}} - e^{-\frac{\sigma_\theta^2+\sigma_\phi^2}{2}})\cos\phi\sin\phi \end{pmatrix} \\ &= \begin{pmatrix} 0 \\ -r(e^{-\frac{\sigma_\phi^2}{2}} - e^{-\frac{\sigma_\theta^2+\sigma_\phi^2}{2}})\cos\phi\sin\phi \end{pmatrix} \end{aligned}$$

So there would be a bias in the second term.

And for virtual measurement $y = \mathbf{h}^*\mathbf{x} + v$
where $\mathbf{h}^* = [\sin\theta, \cos\theta]$, the noise

$$v = r - \mathbf{h}^*\mathbf{x} = r - (x_1\sin\theta - x_2\cos\theta)$$

so that the mean of the noise

$$E[v] = (1 - e^{-\frac{\sigma_\theta^2}{2}})r$$

which means there is bias in mean value, where in the 3D case

$$E[v] = r(1 - e^{-\frac{\sigma_\phi^2}{2}}\sin^2\phi - e^{-\frac{\sigma_\theta^2+\sigma_\phi^2}{2}}\cos^2\phi)$$

Following the same logic and process, we get that in Case II (3D)

$$E[\mathbf{v}] = \begin{pmatrix} 0 \\ -r(e^{-\frac{\sigma_\phi^2}{2}} - e^{-\frac{\sigma_\theta^2+\sigma_\phi^2}{2}})\cos\phi\sin\phi \\ r(1 - e^{-\frac{\sigma_\phi^2}{2}}\sin^2\phi - e^{-\frac{\sigma_\theta^2+\sigma_\phi^2}{2}}\cos^2\phi) \end{pmatrix}$$

In Case III (3D),

$$E[\mathbf{v}] = \begin{pmatrix} 0 \\ -r(e^{-\frac{\sigma_\phi^2}{2}} - e^{-\frac{\sigma_\theta^2+\sigma_\phi^2}{2}})\cos\phi\sin\phi \\ \dot{\theta}(e^{-\frac{\sigma_\theta^2}{2}} - e^{-\frac{\sigma_\phi^2}{2}})r\sin^2\phi + (e^{-\frac{\sigma_\theta^2}{2}} - e^{-\frac{\sigma_\theta^2+\sigma_\phi^2}{2}})r\cos^2\phi) \\ (e^{-\frac{\sigma_\phi^2}{2}} - e^{-\frac{\sigma_\theta^2+\sigma_\phi^2}{2}})(u_3\cos\phi - \dot{\phi}r\sin^2\phi) \end{pmatrix}$$

In Case IV (3D),

$$E[\mathbf{v}] = \begin{pmatrix} 0 \\ -r(e^{-\frac{\sigma_\phi^2}{2}} - e^{-\frac{\sigma_\theta^2+\sigma_\phi^2}{2}})\cos\phi\sin\phi \\ (e^{-\frac{\sigma_\phi^2}{2}} - e^{-\frac{\sigma_\theta^2+\sigma_\phi^2}{2}})(\tau u_3 \sin\phi - r\sin^2\phi) \end{pmatrix}$$

We can see that in each case, the means of noise only shift in with a scale coefficient of $(e^{-\frac{\sigma_\phi^2}{2}} - e^{-\frac{\sigma_\theta^2+\sigma_\phi^2}{2}})$ or $(1 - e^{-\frac{\sigma_\theta^2}{2}})$. When the variances σ_θ and σ_ϕ are small, that coefficient is almost zero. Even when we increase in simulations the actual variances of the bearing measurements to 10° (which is unrealistic based on the performances of current instruments), the mean shift is still in on the scale of 10^{-2}m. It thus remains negligible and does not need to be subtracted explicitly.

Remark

The variance of the noise on the virtual measurements can also be easily approximated. For the bearing and range virtual measurements, one has

$$Var(\mathbf{hx}) \approx \sigma_\theta^2 \, r^2$$

$$Var(\mathbf{h}^*\mathbf{x}) \approx \frac{\sigma_\theta^4}{4} \, r^2$$

$$Cov(\mathbf{h}^*\mathbf{x}, \mathbf{hx}) = 0$$

Since the exact r is not known, when computing R it may be conservatively replaced by a known upper bound r_{max} , or more finely by $r^* = min(r_{measured} + 3\sigma_r \ , \ r_{max})$, where σ_r is the variance of the range measurement noise.

4.5 Extensions

Pinhole camera model

The principle of transforming a nonlinear measurement into a LTV representation is applicable to other contexts, for instance the pinhole camera model (Fig. 2),

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = -\frac{f}{x_3} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

which can be rewritten as LTV constraints on the states (x_1, x_2, x_3) ,

$$\begin{bmatrix} f & 0 & y_1 \\ 0 & f & y_2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \mathbf{Hx} = \mathbf{0}$$

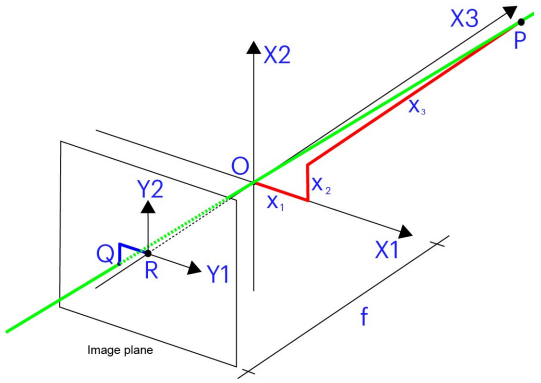

Fig. 3 Pinhole Camera Model

based on the measured y_1 and y_2. If in addition we measure the velocity $\mathbf{u}$ of the camera center and the angular velocity matrix Ω describing the vehicle's rotation, the kinematics model is

$$\dot{\mathbf{x}} = -\mathbf{u} - \Omega\mathbf{x}$$

So we can extend the same LTV Kalman system to estimate the local position of the target shown on the image plane as:

$$\dot{\hat{\mathbf{x}}} = -\mathbf{u} - \Omega\hat{\mathbf{x}} - \mathbf{P}\mathbf{H}^T\mathbf{R}^{-1}\mathbf{H}\hat{\mathbf{x}}$$

with covariance updates

$$\dot{\mathbf{P}} = \mathbf{Q} - \mathbf{P}\mathbf{H}^T\mathbf{R}^{-1}\mathbf{H}\mathbf{P} - \Omega\mathbf{P} + \mathbf{P}\Omega$$

Structure from motion

Structure from motion (Soatto, Frezza and Perona, 1996; Soatto and Perona, 1997, 1998) is a problem in machine vision field dealing with range imaging. The problem is to recover the 3D model of a structure (building, furniture, etc.) from a series of 2D images. The simplified structure from motion problem can be modeled as a combination of SLAM and pin-hole camera. Intuitively, it may be a global version of the pin-hole camera model we introduced previously, where estimations on local positions of the features are replaced by estimations on global positions of features along with global pose and attitude of the camera. So similar with the discussions in the pin-hole model, we can easily write the LTV constraints about any observation as:

$$\begin{bmatrix} f & 0 & y_{i1} \\ 0 & f & y_{i2} \end{bmatrix} \mathbf{T}(\beta)(\mathbf{x}_i - \mathbf{x}_c) = \mathbf{0}$$

where $\mathbf{T}(\beta)$ is the rotation matrix from global coordinates to local coordinates on the camera, and $\mathbf{x}_i, \mathbf{x}_c$ are respectively position of a feature and position of the camera. As in the method we used in the global SLAM discussion, we can have a separate estimation on β and treat it as an input to the LTV Kalman filter as:

$$\dot{\hat{\beta}} = \omega + (\beta_d - \hat{\beta})$$

while kinematics of the estimated states are simply:

$$\dot{\mathbf{x}}_i = \mathbf{0}$$

$$\dot{\mathbf{x}}_c = \mathbf{u}$$

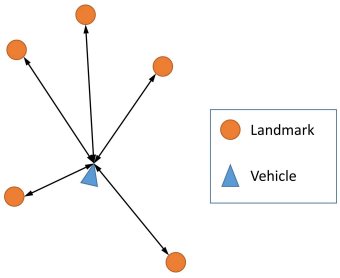

Fig. 4 Graph model of EKF-SLAM: all states including both the landmarks and the vehicle are coupled together

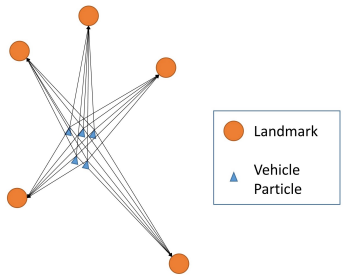

Fig. 5 Graph model of Fast-SLAM: states of landmarks are fully decoupled conditioned on each particle of vehicle states

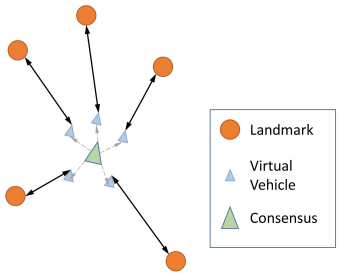

Fig. 6 Graph model of SLAM-DUNK: states of each landmark are only coupled with the corresponding virtual vehicle, and consensus of virtual vehicles as maximization of likelihood is used as best estimate

5 Decoupled Unlinearized Networked Kalman-filter (SLAM-DUNK)

One of the main problems for the proposed full LTV Kalman filter in global coordinates is computation complexity. As shown in Fig.4, since all landmarks are coupled to each other by the vehicle's state, we will have to deal with a full covariance matrix, which requires $O(n^2)$ storage and $O(n^3)$ computation in each step, where n is the number of landmarks. This key limitation restrains the algorithm from being applied to large-scale environment models that could contain tens of thousands of features easily.

While actually the SLAM problem exhibits important conditional independence: that is, conditioned on the vehicle's states of path, all landmarks are decoupled and independent of each other, as suggested in Montemerlo et al. (2002) and Montemerlo and Thrun (2007). In other words, if we feed the vehicle states estimated by other methods into the filters as prior information, we can decoupled the full LTV Kalman filter into n independent location estimation problems, one for each landmark. For example, in Montemerlo et al. (2002) and Montemerlo and Thrun (2007) factorized the SLAM problem into a graph model like Fig.5, where they use M particle filters to update the states of vehicle, and each particle of vehicle is connected to n independent EKF estimators, so that there would be nM filters in total, which means $O(nM)$ computation complexity.

We are proposing a novel algorithm that can decouple the covariance between landmarks into smaller independent estimators and requires less computation even comparing to FastSLAM. Instead of dealing with all measurements and landmarks as one whole state vector to estimate, which is what the full filters like EKF are doing, we want to process information from each single measurement independently with one specific virtual vehicle. Then we establish a consensus summarizing all the information and feedback to the individual observers. Graph model of the proposed algorithm is shown in Fig. 6, and detailed algorithm design is introduced below.

Distributed Sensing

For each landmark $\mathbf{x}_i$, we assign a virtual vehicle $\mathbf{x}_{vi}$ exclusively to process any information generated from that landmark. Using the case where we have both range and bearing measurements for example, the linear constraints between landmark $\mathbf{x}_i$ and virtual vehicle $\mathbf{x}_{vi}$ would be:

$$\mathbf{y}_{i1} = \mathbf{H}_{i1} \begin{bmatrix} \mathbf{x}_i \\ \mathbf{x}_{vi} \end{bmatrix}$$

where

$$\mathbf{y}_{i1} = \begin{bmatrix} 0 \\ r_i \end{bmatrix}$$

and

$$\mathbf{H}_{i1} = \begin{bmatrix} \sin\theta_i & -\cos\theta_i & sin\theta_i & \cos\theta_i \\ \cos\theta_i & \sin\theta_i & -\cos\theta_i & -\sin\theta_i \end{bmatrix} \mathbf{T}(\beta)$$

Such constraint is similar with the one we discussed before, just with additional rotation term because of global coordinates. In that case, each landmark $\mathbf{x}_i$ is coupled with a virtual vehicle $\mathbf{x}_{vi}$ exclusively. And using LTV Kalman filter for each pair, information from observation of any single landmark gets conveyed to the virtual vehicle layer.

Consensus among virtual vehicles

Then in the layer of virtual vehicles, we summarize information from all observations to get a consensus, and use that consensus to guide all virtual vehicles to follow, which makes to virtual vehicle layer a "leader-follower" network. The consensus $\mathbf{x}_{vc}$ is achieve from a weighted average among all virtual vehicles whose corresponding landmarks are observed right now as:

$$\mathbf{x}_{vc} = \frac{\sum_{i \in O} \frac{x_{vi}}{\sigma_{vi}^2}}{\sum_{i \in O} \frac{1}{\sigma_{vi}^2}}$$

The weighted average above is the direct maximization of likelihood summarizing information from all observations. We can simultaneously feed the consensus result to the whole network as a leader for all $\mathbf{x}_{vi}$'s, which can be included into the first level LTV Kalman filter that each virtual vehicle is involved with as an extra observation where

$$\mathbf{y}_{i2} = \mathbf{x}_{vc}$$

$$\mathbf{H}_{i2} = [\mathbf{0} \quad I]$$

Complete Algorithm

So to summarize, the algorithm we propose here is composed of two levels of computation: the first level uses separate LTV Kalman filters for each single pair of landmark and virtual vehicle, including both the measurement and the following behavior towards the consensus for the virtual vehicle. The second level is to gather information from all virtual vehicles who have their corresponding landmarks under

observation. The consensus $\mathbf{x}_{vc}$ is the best estimation from the weighted average that maximizes the likelihood.
For all landmarks $\mathbf{x}_i$ and virtual vehicles $\mathbf{x}_{vi}$:

$$\mathbf{y}_i = \begin{bmatrix} \mathbf{y_{i1}} \\ \mathbf{y_{i2}} \end{bmatrix} \quad \text{and} \quad \mathbf{H}_i = \begin{bmatrix} \mathbf{H}_{i1} \\ \mathbf{H}_{i2} \end{bmatrix}$$

And the LTV Kalman filter for each virtual vehicle is:

$$\begin{bmatrix} \dot{\hat{\mathbf{x}}}_i \\ \dot{\hat{\mathbf{x}}}_{vi} \end{bmatrix} = \begin{bmatrix} 0 \\ \mathbf{u} \end{bmatrix} + \mathbf{P}_i \mathbf{H}_i^T \mathbf{R}^{-1} (\mathbf{y}_i - \mathbf{H}_i \begin{bmatrix} \hat{\mathbf{x}}_i \\ \hat{\mathbf{x}}_{vi} \end{bmatrix})$$

$$\dot{\mathbf{P}}_i = \mathbf{Q} - \mathbf{P}_i \mathbf{H}_i^T \mathbf{R}^{-1} \mathbf{H}_i \mathbf{P}_i$$

Remarks
Under certain situations, there would be some special cases for the proposed algorithm. When the landmark i is not observed by the vehicle right now, the observation parts of the first level estimator would be dropped, that means $\mathbf{y}_{i1}$ and $\mathbf{H}_{i1}$ would not be included. When the vehicle sees no landmark at any moment, since there would be no way to achieve $\mathbf{x}_{vc}$ from weight average, $\mathbf{y}_{i2}$ and $\mathbf{H}_{i2}$ would be dropped.
When the vehicle see a new landmark for the first time, the corresponding virtual vehicle is initialized at the location of current best estimation, which can be computed from weight average among all virtual vehicles as:

$$\mathbf{x}_{vi0} = \frac{\sum_{j \in all} \frac{x_{vj}}{\sigma_{vj}^2}}{\sum_{j \in all} \frac{1}{\sigma_{vj}^2}}$$

Furthermore, data association to match the observed landmarks with the ones in the memory can be simply carried by matching the measurements with the saved pairs of landmarks and virtual vehicles. Notice that since for each pair of landmark and virtual vehicle, we design a specific LTV Kalman filter specifically. That means we would have n filters in total, where n is the total number of landmarks. For each filter, it would have two states to estimate. So the total computation complexity would be $O(n)$, which would be comparable to FastSLAM with only two particles, while FastSLAM with only two particles would sacrifice hugely on performance. Moreover, for filters whose landmarks that are not observed, they only have the behavior of following, so the computation is even lighter.
The whole idea of the proposed algorithm is to break the unified full LTV Kalman filter containing both the landmarks and the vehicle states into n small estimators to get the best estimation locally and one more maximization of likelihood to achieve the best estimation on consensus. We successfully decouple the landmarks while keep the whole method theoretically sound. And keep in mind that each single estimator still follows the same structure as the full ones proposed before, so contraction analysis as same as the full LTV Kalman filters can be exploited to ensure the estimations would finally converge to the true states.
Despite the authors' lack of enthusiasm for acronyms, it is hard to resist calling

the final algorithm SLAM-DUNK [1], for simultaneous location and mapping using distributed unlinearized networked Kalman-filtering.

6 Experiments

Experiments for 2D landmarks estimation
We experiment the 2D version of our cases with simulations in Matlab. As shown in `https://vimeo.com/channels/910603`, in the simulations, we have three landmarks. The diameter of each landmark is $d = 2$m. We have run simulations on all five cases. The noise signals that we use in the simulations are: standard variance for zero-mean Gaussian noise of θ is 2°; standard variance for noise of $\dot{\theta}$ is 5°/s; standard variance for measurement noise of r is $2m$; and standard variance for noise of α is 0.5°. In `https://vimeo.com/channels/910603`, the green lines

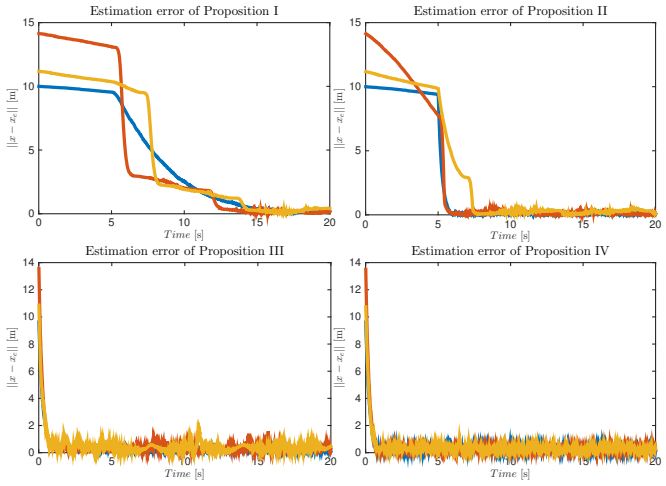
Estimation error of Proposition I
Estimation error of Proposition II
Estimation error of Proposition III
Estimation error of Proposition IV
Time [s]

Fig. 7 Estimation error for Case I, II, III, IV for 2D estimation

indicate the trajectories of estimations of the each case. The blue lines are the movement trajectory of the vehicle. As shown in `https://vimeo.com/channels/910603`, trajectories of estimations from Case II, III and IV are smoother and converge faster than original Case I and Case V. This is because the trajectory exploits additional information. In particular, for Case II and IV, since the "time-to-contact" measurement and radial distance measurement both contains information on the ra-

[1] We thank Geoffrey Hinton for the suggestion.

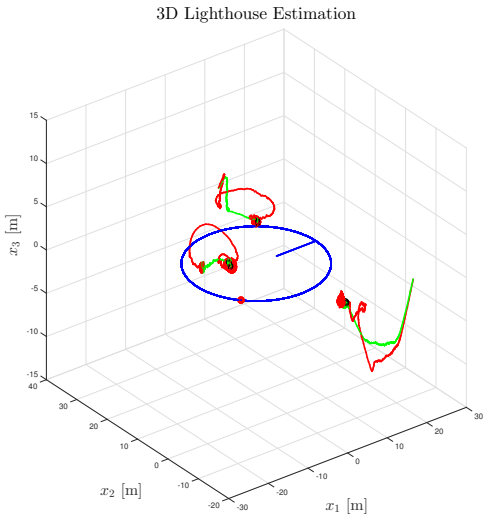

Fig. 8 3D landmarks estimation in Case I and III

dial direction, they converge to the true position directly, without waiting for the vehicle movement to bring in extra information.

Next, we analyze the estimation errors $||\mathbf{x} - \hat{\mathbf{x}}||$ in Case I, II, III and IV for all three landmarks in Fig. 7. The figure shows that the errors decay faster in Case II, III and IV. The difference is that Case III needs to wait for the movement of the vehicle to provide more information about $\dot{\theta}$, yet Case II and IV contract much faster with exponential rates because of range related measurements. Compared to Case II, Case IV is less smooth, as expected, because the time-to-contact measurement itself is an approximation and may be disturbed when $\dot{r}$ is close to zero.

Experiments for 3D landmarks estimation

We also have simulation results for Case I and Case III in 3D settings. Here we also have three lighthouses with different locations. The results shown in Fig. 8 suggest that our algorithm is capable of estimating landmarks positions accurately in 3D space with bearing angle for both yaw and pitch. Animations of all simulation results are provided at `https://vimeo.com/channels/910603`.

Experiment for Victoria Park landmarks estimation

We applied our algorithm to Sydney Victoria Park dataset, a popular dataset in the SLAM community. The vehicle path around the park is about 30 min, covering over 3.5 km. Landmarks in the park are mostly trees. Estimation results are compared with intermittent GPS information as ground truth to validate the states of the filters as shown in Fig. 9. Our estimated track compares favorably to benchmark result of Kim, Sakthivel and Chung (2008), which highlights the consistency of our al-

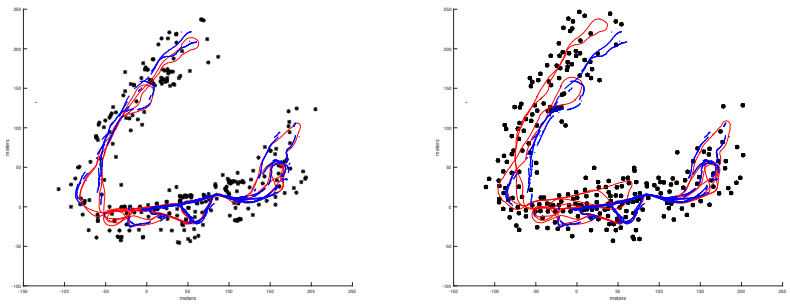

Fig. 9 On the left is the path and landmarks estimation of our algorithm and on the right is the result from Unscented Fast SLAM. The thick blue path is the GPS data and the solid red path is the estimated path; the black asterisks are the estimated positions of the landmark.

gorithm in large scale applications. Simulation result of the Victoria Park dataset is provided at `https://vimeo.com/136219156`.

7 Concluding Remarks

In this paper, we propose using the combination of LTV Kalman filter and contraction tools to solve the problem of simultaneous mapping and localization (SLAM). By exploiting the virtual measurements, the LTV Kalman observer does not suffer from errors brought by the linearization process in the EKF SLAM, which makes the solution global and exact. Convergence rates can be quantified using contraction analysis. The application cases utilize different kinds of sensor information that range from traditional bearing measurements and range measurements to novel ones like optical flows and time-to-contact measurements. They can solve SLAM problems in both 2D and 3D scenarios. Note that

- bounding of the covariance matrix $\mathbf{P}$ may be done analytically based on the observability Grammian (Bryson, 1975; Lohmiller and Slotine, 2013a).
- our approach is particularly suitable for exploiting the recent availability of vision sensors at very low cost, rather than relying on range sensors like lidars.
- in the Victoria Park benchmark dataset, features are mostly trees in the park. As a result some regions have dense landmarks, while others have sparse landmarks. Landmarks in dense areas and landmarks with high uncertainty provides less information for the updates on the states. Thus, incorporating feature selection to use landmarks with richer information could reduce computation workload like suggested in Mu (2013) . Such active sensing could be achieved as in (Dickmanns, 1998, 2007; Slotine and Lohmiller, 2001) by exploiting the fact that the

posterior covariance matrix can be computed before taking any specific measurement. More generally, path planning may also be adjusted according to a desired exploration/exploitation trade-off (Vergassola, Villermaux and Shraiman, 2007; Schwager, Rus and Slotine, 2009; Mu, 2013).
- it may be interesting to consider whether similar representations may also be used in biological navigation, e.g. in the context of place cells or grid cells (Moser, Kropff and Moser, 2008) or sensing itself (Gollisch and Meister, 2010).

References

Bailey T, Nieto J, Guivant J, Stevens M and Nebot E (2006) Consistency of the EKF-SLAM algorithm. In: *Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on*. IEEE, pp. 3562–3568.

Browning NA (2012) A neural circuit for robust time-to-contact estimation based on primate MST. *Neural Comput.* 24(11): 2946–2963.

Bryson AE (1975) *Applied optimal control: optimization, estimation and control*. CRC Press.

Cheeseman P, Smith R and Self M (1987) A stochastic map for uncertain spatial relationships. In: *4th International Symposium on Robotic Research*. pp. 467–474.

Clady X, Clercq C, Ieng SH, Houseini F, Randazzo M, Natale L, Bartolozzi C and Benosman R (2015) Asynchronous visual event-based time-to-contact. *Neuromorphic Engineering Systems and Applications* : 51.

Del Moral P and Guionnet A (1999) Central limit theorem for nonlinear filtering and interacting particle systems. *Annals of Applied Probability* : 275–297.

Dellaert F and Kaess M (2006) Square Root SAM: Simultaneous localization and mapping via square root information smoothing. *The International Journal of Robotics Research* 25(12): 1181–1203.

Demidovich BP (1961) Dissipativty of a system of nonlinear differential equations. *Ser. Mat. Mekh.* .

Dickmanns ED (1998) Vehicles capable of dynamic vision: a new breed of technical beings? *Artificial Intelligence* 103(1): 49–76.

Dickmanns ED (2007) *Dynamic vision for perception and control of motion*. Springer Science & Business Media.

Dogar MR, Hemrajani V, Leeds D, Kane B and Srinivasa S (2010) Proprioceptive localization for mobile manipulators .

Doucet A, De Freitas N, Murphy K and Russell S (2000) Rao-Blackwellised particle filtering for dynamic Bayesian networks. In: *Proceedings of the Sixteenth conference on Uncertainty in artificial intelligence*. Morgan Kaufmann Publishers Inc., pp. 176–183.

Duckett T, Marsland S and Shapiro J (2002) Fast, on-line learning of globally consistent maps. *Autonomous Robots* 12(3): 287–300.

Folkesson J and Christensen H (2004) Graphical SLAM a self-correcting map. In: *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, volume 1. IEEE, pp. 383–390.

Gollisch T and Meister M (2010) Eye smarter than scientists believed: neural computations in circuits of the retina. *Neuron* 65(2): 150–164.

Grave I and Tang Y (2015) A new observer for perspective vision systems under noisy measurements. *Automatic Control, IEEE Transactions on* 60(2): 503–508.

Grisetti G, Kummerle R, Stachniss C and Burgard W (2010) A tutorial on graph-based SLAM. *Intelligent Transportation Systems Magazine, IEEE* 2(4): 31–43.

Hartman P (1964) *Ordinary Differential Equations*. John Wiley & Sons.

Horn BK, Fang Y and Masaki I (2007) Time to contact relative to a planar surface. In: *IEEE intelligent vehicles symposium*. pp. 68–74.

Huang S and Dissanayake G (2007) Convergence and consistency analysis for extended Kalman filter based SLAM. *Robotics, IEEE Transactions on* 23(5): 1036–1049.

Javdani S, Klingensmith M, Bagnell JA, Pollard NS and Srinivasa SS (2013) Efficient touch based localization through submodularity. In: *Robotics and Automation (ICRA), 2013 IEEE International Conference on*. IEEE, pp. 1828–1835.

Jazwinski AH (2007) *Stochastic processes and filtering theory*. Courier Corporation.

Kalman RE (1960) A new approach to linear filtering and prediction problems. *Journal of Fluids Engineering* 82(1): 35–45.

Kim C, Sakthivel R and Chung WK (2008) Unscented FastSLAM: a robust and efficient solution to the SLAM problem. *Robotics, IEEE Transactions on* 24(4): 808–820.

Klingensmith M, Srinivasa S and Kaess M (2016) Articulated robot motion for simultaneous localization and mapping (arm-slam). *IEEE Robotics and Automation - Letters* .

Konolige K (2004) Large-scale map-making. In: *Proceedings of the National Conference on Artificial Intelligence*. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, pp. 457–463.

Lewis D (1949) Metric properties of differential equations. *American Journal of Mathematics* : 294–312.

Lohmiller W and Slotine JJ (1998) On contraction analysis for non-linear systems. *Automatica* 34(6): 683–696.

Lohmiller W and Slotine JJ (2008) Shaping state-dependent convergence rates in nonlinear control system design. In: *AIAA Guidance, Navigation, and Control Conference*.

Lohmiller W and Slotine JJ (2013a) Contraction analysis of nonlinear Hamiltonian systems. In: *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*. IEEE, pp. 6586–6592.

Lohmiller W and Slotine JJ (2013b) Exact decomposition and contraction analysis of nonlinear hamiltonian systems. In: *AIAA Guidance, Navigation, and Control Conference*.

Matthies L and Shafer SA (1987) Error modeling in stereo navigation. *Robotics and Automation, IEEE Journal of* 3(3): 239–248.

Montemerlo M and Thrun S (2006) Large-scale robotic 3-D mapping of urban structures. In: *Experimental Robotics IX*. Springer, pp. 141–150.

Montemerlo M and Thrun S (2007) Fastslam 2.0. *FastSLAM: A Scalable Method for the Simultaneous Localization and Mapping Problem in Robotics* : 63–90.

Montemerlo M, Thrun S, Koller D, Wegbreit B et al. (2002) FastSLAM: A factored solution to the simultaneous localization and mapping problem. In: *AAAI/IAAI*. pp. 593–598.

Moser EI, Kropff E and Moser MB (2008) Place cells, grid cells, and the brain's spatial representation system. *Annu. Rev. Neurosci.* 31: 69–89.

Moutarlier P and Chatila R (1989) Stochastic multisensory data fusion for mobile robot location and environment modeling. In: *International Symposium of Robotics Research*.

Moutarlier P and Chatila R (1990) An experimental system for incremental environment modelling by an autonomous mobile robot. In: *Experimental Robotics I*. Springer, pp. 327–346.

Mu B (2013) *Value of information based distributed inference and planning*. PhD Thesis, Massachusetts Institute of Technology.

Schwager M, Rus D and Slotine JJ (2009) Decentralized, adaptive coverage control for networked robots. *The International Journal of Robotics Research* 28(3): 357–375.

Slotine JJ and Lohmiller W (2001) Modularity, evolution, and the binding problem: a view from stability theory. *Neural networks* 14(2): 137–145.

Smith R, Self M and Cheeseman P (1990) Estimating uncertain spatial relationships in robotics. In: *Autonomous robot vehicles*. Springer, pp. 167–193.

Soatto S, Frezza R and Perona P (1996) Motion estimation via dynamic vision. *IEEE Transactions on Automatic Control* 41(3): 393–413.

Soatto S and Perona P (1997) Recursive 3-d visual motion estimation using subspace constraints. *International Journal of Computer Vision* 22(3): 235–259.

Soatto S and Perona P (1998) Reducing structure from motion: a general framework for dynamic vision. 1. modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20(9): 933–942.

Thrun S and Montemerlo M (2006) The graph SLAM algorithm with applications to large-scale mapping of urban structures. *The International Journal of Robotics Research* 25(5-6): 403–429.

Vergassola M, Villermaux E and Shraiman BI (2007) Infotaxis as a strategy for searching without gradients. *Nature* 445(7126): 406–409.